

Performance of Opportunistic Epidemic Routing on Edge-Markovian Dynamic Graphs

John Whitbeck, Vania Conan, and Marcelo Dias de Amorim

Abstract

Connectivity patterns in intermittently-connected mobile networks (ICMN) can be modeled as edge-Markovian dynamic graphs. We propose a new model for epidemic propagation on such graphs and calculate a closed-form expression that links the best achievable delivery ratio to common ICMN parameters such as message size, maximum tolerated delay, and link lifetime. These theoretical results are compared to those obtained by replaying a real-life contact trace.

Index Terms

Intermittently-connected Mobile Networks, Network Modeling, Markovian Random Graphs, Epidemic Routing

I. INTRODUCTION

Intermittently connected mobile networks (ICMN) emerge from the social processes that bring mobile devices into contact. Due to high node mobility and frequent lack of end-to-end connectivity in such networks, message transport is usually handled in a store-and-forward fashion by delay/disruptive-tolerant network (DTN) routing protocols [1].

The topology of a real-life network of mobile devices evolves over time as links come up and down. A network's connectivity graph is defined by associating each node to a vertex and adding

A poster of this work was presented at the *ACM SIGCOMM Workshop on Networking, Systems, Applications on Mobile Handhelds (Mobiheld 2009)*. This version is more detailed and contains many more results.

This work has been partially supported by the ANR project Crowd under contract ANR-08-VERS-006.

John Whitbeck is with both LIP6/CNRS – UPMC Sorbonne Universités and Thalès Communications, France. Email: john.whitbeck@lip6.fr. Vania Conan is with Thalès Communications, France. Email: vania.conan@fr.thalesgroup.com. Marcelo Dias de Amorim is with LIP6/CNRS – UPMC Sorbonne Universités, France. Email: marcelo.amorim@lip6.fr.

an edge between any pair of nodes that are currently in contact (i.e., within transmission range of each other). Successive snapshots of the evolving connectivity graph yield a *dynamic graph*, i.e., a time-indexed sequence of static connectivity graphs. Their theoretical study is therefore important for understanding the underlying network dynamics.

In this paper, we propose a new Markovian model for flooding on edge-Markovian dynamic graphs [2]. Unlike previous work on asymptotic behavior [2], our approach assumes source-destination pairs for messages, a finite number of nodes, as well as finite link capacities and message sizes. Our main contribution is a closed-form expression of the bundle¹ delivery probability as a function of bundle size, maximum tolerated delay, and the dynamics of the underlying edge-Markovian dynamic graph. Using this model, we show that the message delivery ratio increases for smaller bundles, but that the achieved gain is bounded and only significant when the constraints on message delivery delay are tight. Finally, we compare our model's predictions to results from a real-life connectivity trace obtained in a rollerblading tour.

In Section II, we briefly describe the edge-Markovian dynamic graph model. We then calculate the delivery ratio for epidemic routing in Section III, before discussing the impact of bundle size on delivery ratio in Section IV. We then compare these theoretical insights to results from a real data set, the Rollernet experiment, in Section V.

II. EDGE-MARKOVIAN DYNAMIC GRAPHS

In the rest of this paper, in both the theoretical and experimental parts, we represent the ad hoc network formed by N mobile nodes as a connectivity graph that evolves in discrete time. Depending on the context, we will use the terms vertex (resp. edge) and node (resp. link) interchangeably. The time step τ is equal to the shortest contact or inter-contact time. In a real-life trace, τ is equal to the neighborhood scanning sampling period. Edges come up or down at the beginning of each time step, but the topology then remains static until the next time step.

Previous work on dynamic graphs focused on graphs with increasing numbers of vertices or edges [3], but did not account for node mobility and/or link instability. More recently, Chaintreau *et al.* used simple sequences of uniform random graphs for modeling random temporal graphs in order to analyze the diameter of opportunistic mobile networks [4]. Pellegrini *et al.* explored

¹Bundles are *message aggregates*. They can contain anything from a message fragment to several messages [1].

the notion of *connectivity over time* but this approach loses all information about the order in which contact opportunities appear [5]. Unfortunately, none of these models capture the strong correlation between successive connectivity graphs.

Edge-Markovian dynamic random graphs were first introduced by Clementi *et al.* as a generalization of time-independent dynamic random graphs to capture the strong dependence between the existence of an edge at a given time step and its existence at the previous time step [2]. While such a model may be used to study a wide variety of dynamic graphs, in this paper, we will focus on its application to ICMNs.

Dynamic random graph based models, including the one in this paper, have an exponential (or geometric) inter-contact time distribution. In real-life datasets this may not always be the case. Indeed, when the underlying social dynamics are strong, the inter-contact distribution follows a power law [6]. However, in different scenarios, it may follow an exponential law [7]. Interestingly, the inter-contact distribution of any mobility model in a bounded domain necessarily exhibits an exponential cutoff [8].

In an edge-Markovian dynamic graph with N vertices, each edge is considered independently and can be in one of two states: either \uparrow or \downarrow . Let p_\uparrow (resp. p_\downarrow) be the probability of transitioning to the \uparrow (resp. \downarrow) state. The transition matrix for each edge is therefore

$$M = \left(\begin{array}{c|cc} & \downarrow & \uparrow \\ \hline \downarrow & 1 - p_\uparrow & p_\uparrow \\ \uparrow & p_\downarrow & 1 - p_\downarrow \end{array} \right). \quad (1)$$

The contact (T_\uparrow) and inter-contact (T_\downarrow) times are distributed geometrically and their expected values are $E(T_\uparrow) = \frac{\tau}{p_\downarrow}$ and $E(T_\downarrow) = \frac{\tau}{p_\uparrow}$. Indeed, the number of time steps required to leave the \uparrow (resp. \downarrow) state is the number of trials needed to get one success in a Bernoulli process with probability p_\downarrow (resp. p_\uparrow). Let π_\uparrow (resp. π_\downarrow) be the stationary probability of being in state \uparrow (resp. \downarrow). We have $\pi_\uparrow = \frac{p_\uparrow}{p_\uparrow + p_\downarrow}$ and $\pi_\downarrow = \frac{p_\downarrow}{p_\uparrow + p_\downarrow}$. Finally, the average node degree is $(N - 1)\pi_\uparrow$.

III. TUNING MESSAGE SIZE TO MEET DELAY CONSTRAINTS

A. Preliminaries

We assume that, when up, all links have equal capacity ϕ and thus can transport the same quantity $\phi\tau$ of information during one time step. We refer to $\phi\tau$ as the *link size*. Small values of τ therefore mean that the network topology's characteristic evolution time is short and thus only

small amounts of information may be transmitted over a link during one time step. We define the *bundle size* as numerically proportional to τ : $\alpha\phi\tau$. By abuse of language, however, we will simply refer to α as the bundle size. For example, a bundle of size 2 ($\alpha = 2$) is only able to traverse links that last for more than 2 time steps, whereas a bundle of size 0.5 is able to traverse two links during each time step. Furthermore, each bundle can only tolerate a certain maximum delay. We note d the maximum number of time steps, beyond which a delivery is considered to have failed. By abuse of language, we will simply refer to d as the maximum delay.

Epidemic routing was one of the first methods proposed for dealing with intermittent connectivity in mobile ad hoc networks [9]. Each message is flooded into the entire network. Upon meeting, two nodes first exchange message vectors describing which messages they currently hold, before requesting from one another copies of the messages they do not yet have. Following the epidemic analogy, a node is said to be *infected* by a message upon receiving a copy of it [10]. Epidemic routing is particularly useful for theoretical purposes, since its delivery ratio is also that of the optimal single-copy time-space routing protocol.

Our goal is to calculate the delivery probability of a bundle using epidemic routing. To be successful, delivery has to occur without exceeding the maximum allowed delay. To this end, we introduce a new Markovian model for epidemic propagation on the edge-Markovian dynamic graph of the previous Section.

For the sake of simplicity, the model will first be described for $\alpha = 1$. In Sections III-C and III-D, we will respectively describe how to adapt the previous model when the bundles are smaller ($\alpha < 1$) and larger ($\alpha > 1$) than the link size.

B. Bundles fit in a time slot ($\alpha = 1$)

Source a wishes to transmit a bundle to destination b using epidemic routing. Edges change states at the beginning of each time step. During one time step, an infected node infects all of its direct uninfected neighbors, and only those, since the bundle size is 1 and bundles can therefore only perform one hop per time step. Let V be the set of the nodes in the network. After k time steps, nodes other than b fall in one of three disjoint sets:

- Those that have *just* been infected: J_k .
- Those that have been infected at time step $k - 1$ or before: I_k .
- Those that have not yet been infected: $S_k = V \setminus (I_k \cup J_k \cup \{b\})$.

This distinction is necessary to determine who can be infected at time step $k + 1$. Indeed, if a node belongs to I_k , then all its neighbors at the end of time step k are in $I_k \cup J_k$. It can only infect new nodes if an edge to a clean node in S_k comes up at time step $k + 1$. However, a node in J_k may have edges to some clean neighbors in S_k which may become infected at time step $k + 1$ if the edge remains up.

In this paper, we are only interested in the probability that b receives a copy in at most d time steps. In this case, the only information necessary to characterize the state of the epidemic is the number of nodes i and j in I_k and J_k , respectively. The delivery ratio can be obtained as the absorption probability of the Markov chain described hereafter.

States. The epidemic can be described as a Markov chain on the following $2 + \frac{N(N-1)}{2}$ states:

- *Init*: The initial state in which only the origin a is infected. This state is transient.
- *Succ*: The destination b has been infected. This state is absorbing.
- States (i, j) for $1 \leq i \leq N - 1$ and $0 \leq j \leq N - 1 - i$. These are also transient.

Primitives. The transition probabilities are functions of the following primitives. Given two sets of nodes U and W , if each node of U can infect each other node in W with probability p , we define the probability that m nodes in W will be infected:

$$P_{inf}(m, p, |U|, |W|) = \text{pdf}_B(m, 1 - (1 - p)^{|U|}, |W|), \quad (2)$$

where $\text{pdf}_B(m, p, n)$ is the probability density function of a binomial distribution of n independent events with probability p .

A node that has just been infected (i.e., $\in J_k$) can contaminate the destination the following round with probability π_\uparrow , while nodes that have been infected for two or more time steps (i.e., $\in I_k$) can do so with probability p_\uparrow . If $|I_k| = i$ and $|J_k| = j$, then the probability of infecting the destination b during the next time step is:

$$P_{succ}(i, j) = 1 - \pi_\downarrow^j (1 - p_\uparrow)^i. \quad (3)$$

Transition Probabilities. The state *Succ* is absorbing. Any transitions from the *Init* state, can be calculated as transitions from a $(0, 1)$ state. A state (i, j) can transition to either state *Succ* with probability $P_{succ}(i, j)$ or to another state $(i + j, j')$ with probability:

$$(1 - P_{succ}(i, j)) \sum_{m=0}^{j'} (P_{inf}(m, \pi_\uparrow, j, N - 1 - i - j) P_{inf}(j' - m, p_\uparrow, i, N - 1 - i - j - m)). \quad (4)$$

Delivery Ratio. Let \mathbf{T} be the Markov chain's matrix of transition probabilities, \mathbf{i} the initial state vector and \mathbf{s} the state vector with coefficient 1 for state *Succ* and 0 for all others. Therefore, the delivery probability (i.e., the probability of being in state *Succ*) after d time steps is $P_{deliv}(d, \alpha = 1) = \mathbf{i}\mathbf{T}^d\mathbf{s}$.

For example, let us consider a network with 3 mobile nodes. The Markov chain describing an epidemic propagation on its associated edge-Markovian dynamic graph has 5 states: *Init* = (0, 1), (1, 0), (1, 1), (2, 0), and *Succ*. Here $\mathbf{s} = [00001]^T$ and its initial vector state vector is $\mathbf{i} = [10000]$. Its matrix of transition probabilities is:

$$\mathbf{T} = \begin{pmatrix} & \text{Init} & (1,0) & (1,1) & (2,0) & \text{Succ} \\ \text{Init} & 0 & \pi_{\downarrow}^2 & \pi_{\downarrow}\pi_{\uparrow} & 0 & \pi_{\uparrow} \\ (1,0) & 0 & (1-p_{\uparrow})^2 & (1-p_{\uparrow})p_{\uparrow} & 0 & p_{\uparrow} \\ (1,1) & 0 & 0 & 0 & \pi_{\downarrow}(1-p_{\uparrow}) & 1-\pi_{\downarrow}(1-p_{\uparrow}) \\ (2,0) & 0 & 0 & 0 & (1-p_{\uparrow})^2 & 1-(1-p_{\uparrow})^2 \\ \text{Succ} & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (5)$$

C. Bundles smaller than link size ($\alpha < 1$)

When the bundle size is smaller than the link size, bundles may perform up to $\lfloor \frac{1}{\alpha} \rfloor$ hops during one time step. Recall that the network topology instantly changes at the beginning of each time step, before the *first* hop. After that, the remaining hops happen on the same *static* network topology. To take this into account, we define a *static* propagation matrix \mathbf{R} using the same states as previously but tweaking the transition probabilities. In a static topology no new links can come up, hence $P_{succ}^{static}(i, j) = 1 - \pi_{\downarrow}^j$ and the transition from state (i, j) to $(i+j, j')$ happens with probability $(1 - P_{succ}^{static}(i, j)) P_{inf}(j', \pi_{\uparrow}, j, N-1-i-j)$. Finally, the delivery ratio (i.e., the probability of being in *Succ* after d time steps) is $P_{deliv}(d, \alpha < 1) = \mathbf{i} \left(\mathbf{T} \cdot \mathbf{R}^{\lfloor \frac{1}{\alpha} \rfloor - 1} \right)^d \mathbf{s}$.

D. Bundles larger than link size ($\alpha > 1$)

Bundles larger than the link size can only use links that last longer than $\lceil \alpha \rceil$ time steps. Computing the exact delivery probability in this case requires one to keep track of the number of nodes that will complete reception of the bundle in $1, 2, \dots, \lceil \alpha \rceil$ time steps. This quickly becomes intractable. Instead one can easily calculate upper and lower bounds on the delivery probability by considering successive, non overlapping, intervals of $\lceil \alpha \rceil$ time steps and only the

links that last longer than $\lceil \alpha \rceil$ time steps. The latter will hereafter be referred to as *sufficiently long* links.

The lower bound is obtained by taking into account only the *sufficiently long* links that either exist or come up at the *beginning* of an interval. This ignores links that come up later in the interval and hence underestimates the propagation of the epidemic. More precisely, we replace π_{\uparrow} by $\pi_{\uparrow}(1 - p_{\downarrow})^{\lceil \alpha \rceil - 1}$ and p_{\uparrow} by $p_{\uparrow}p_{\downarrow}^{\lceil \alpha \rceil - 1}$ in Eqs. 3 and 4.

The upper bound is obtained by considering that *any sufficiently long* link that comes up during one interval will allow the *full* bundle to be transmitted over it by the end of the interval. For example, if $\alpha = 2$ and a *sufficiently long* link appears after one time step within the two-time-step interval, then we consider that the whole bundle can be transferred over that link before the next interval. This obviously overestimates the number of infected nodes at each time step. More precisely, we replace π_{\uparrow} by $(\pi_{\uparrow} + \pi_{\downarrow}(1 - (1 - p_{\uparrow})^{\lceil \alpha \rceil - 1}))(1 - p_{\downarrow})^{\lceil \alpha \rceil - 1}$ and p_{\uparrow} by $(1 - (1 - p_{\uparrow})^{\lceil \alpha \rceil})(1 - p_{\downarrow})^{\lceil \alpha \rceil - 1}$ in Eqs. 3 and 4.

If \mathbf{T}_l (resp. \mathbf{T}_u) is the transition matrix obtained for the lower (resp. upper) bound, then the delivery probability after d time steps is bounded by $\mathbf{i}\mathbf{T}_l^{\frac{d}{\lceil \alpha \rceil}}\mathbf{s} \leq P_{deliv}(d, \alpha > 1) \leq \mathbf{i}\mathbf{T}_u^{\frac{d}{\lceil \alpha \rceil}}\mathbf{s}$.

IV. DISCUSSION

A. Influence of bundle size

Bundles larger than the link size see their delivery probability severely degraded, though this is somewhat mitigated by longer maximum delays (Fig. 1). On the other hand, bundles smaller than the link size can make several hops in a single time step. This is a great advantage when the time constraints are particularly tight ($d = 4$ in Fig. 1), but barely has any effect when the time constraints are looser. This also highlights the influence of node mobility. Indeed, since the actual bundle size is proportional to τ (see Section III-A), high node mobility (i.e., small τ) makes the actual link size smaller and thus further constrains possible bundle size.

B. Influence of other parameters

Number of nodes. (Fig. 2a) The delivery probability tends to 1 as N increases. Indeed, for a given source/destination pair, each new node is a new potential relay in the epidemic dissemination and thus helps the delivery ratio.

Topology evolution speed. (Fig. 2b) Faster oscillations between \uparrow and \downarrow states make for a more dynamic network topology. This makes for shorter contact and inter-contact times (Section II) but increases contact opportunities. Small bundles ($\alpha \leq 1$) take advantage of this and their delivery ratio increases as $E(T_{\uparrow})$ decreases. On the other hand, excessive link instability drives the delivery ratio for larger bundles ($\alpha > 1$) to 0, because fewer links last longer than one time step.

Average node degree. (Fig. 2c) Greater connectivity increases the delivery probability. The sharp slope of the curve when $\alpha \leq 1$ is reminiscent of percolation in random graphs when the average node degree hits 1.

Maximum Delay. (Fig. 2d) All else being equal, there is a threshold value beyond which almost all bundles are delivered. This can be linked to the space-time diameter of the underlying topology [4].

V. EVALUATION

The theoretical results from the previous section give us valuable insights into real-life scenarios. Although the edge-Markovian model's diameter is significantly smaller than that of real-work networks due to unwanted small-world properties, it accurately predicts, as we shall see in this section, the relations between delivery ratio, maximum delay and bundle size.

A. Methodology

Wireless connectivity traces involving mobile devices have typically been conducted using periodic Bluetooth scans [6], [11], [12]. In this paper, we chose to study the Rollernet trace [12], which captures the connectivity patterns in a rollerblading tour, because of its very short sampling period. Indeed, the longer the sampling period, the more likely link failures or short contacts will be missed. Furthermore, it becomes difficult to claim that a contact translates into a link that lasts roughly as long as the sampling period (one of our core theoretical assumptions). Therefore, in order to compare theoretical and experimental results, we require traces with very short sampling periods.

Other Bluetooth contact traces were considered, such as the Reality Mining experiment conducted at MIT [11] or the Infocom 2005 traces from the Huggle Project [6]. Unfortunately, none of these had a short enough sampling period (600 and 120 seconds respectively, compared

to 15 seconds for Rollernet). In a sense, the MIT and Infocom traces capture a subset of contact opportunities while Rollernet approaches the evolution of the connectivity graph.

Since the dataset logs contacts between nodes and not link durations, we assumed that two nodes in contact remain so for the entire sampling period. Furthermore, we did not try to extrapolate additional events (e.g., new contact opportunities and link failures) between multiples of the sampling period. As in the Markovian network model described previously, we again assume that all links have equal capacity. The first 3,000 seconds of Rollernet trace were replayed. Every 15 seconds for the first 2,000 seconds, 60 source/destination pairs were randomly selected for a simulation of epidemic routing. The average link lifetime is 26.18 seconds and the average node degree 4.75. Using the expressions from Section II, we derive $p_{\uparrow} = 0.05$ and $p_{\downarrow} = 0.57$.

B. Results

Predicting deliver ratio. Fig. 3 compares the measured delivery ratio in Rollernet to the model's upper and lower bounds. Here we use the p_{\uparrow} and p_{\downarrow} values derived from the trace's average link lifetime and node degree. Due to small-world effects, our model is overly optimistic, particularly for smaller bundle sizes ($\alpha \leq 2$). However it does successfully bound the experimental values for larger bundle sizes.

Smaller bundles increase delivery ratio. In Fig. 4, the delivery ratio is steady and close to 1 before dropping sharply beyond a certain bundle size that depends on the target delay. Due to mobility, more than half of the links last less than 15 seconds. Therefore, bundles of size greater than 1 forgo many contact opportunities. However, longer maximum delays can compensate for this. This mirrors the theoretical results on size, delay, and mobility shown in Fig. 1.

Bounded gain from smaller bundles. In Fig. 4, when the maximum delay is 1 minute, the maximum achievable delivery ratio is 0.95 no matter how small the bundles are. This bound on the gain achieved by smaller bundles appears because they hit the performance limit of epidemic routing. Indeed, the best possible epidemic propagation of a message will, at each time step, infect a whole connected component if at least one of its nodes is infected. A small enough bundle can spread sufficiently quickly to achieve this, and thus even smaller bundles bring no performance gain. The same bounded gain from smaller bundles is visible on Fig. 1 on the $d = 4$ curve.

Tight delays require smaller bundles. The sharp delivery ratio drop in Fig. 4 occurs later for

more relaxed delay constraints. A tight time constraint (less than a couple of minutes for example) forces the use of smaller bundles in order to obtain an acceptable delivery ratio. On the other hand, looser time constraints allow for more flexibility regarding bundle size. It is therefore possible to determine the maximum bundle size for any given target delivery ratio.

VI. CONCLUSION

In this paper, we proposed a new model for epidemic propagation on edge-Markovian dynamic graphs which capture the correlation between successive connectivity graphs. We find a closed-form expression of delivery ratio as a function of bundle size, maximum tolerated delay, and the dynamics of the underlying evolving graph. In particular, we have shown that, given a certain maximum delay and node mobility, bundle size has a major impact on the delivery ratio. Our theoretical insights on the interaction between these parameters are corroborated by experimental results on the Rollernet dataset.

REFERENCES

- [1] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proc. ACM SIGCOMM*, 2003.
- [2] A. E. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri, "Flooding time in edge-Markovian dynamic graphs," in *Proc. ACM PODC*, 2008.
- [3] J. Leskovec, J. Kleinberg, and C. Faloutsos, "Graphs over time: Densification laws, shrinking diameters and possible explanations," in *Proc. ACM SIGKDD*, 2005.
- [4] A. Chaintreau, A. Mtibaa, L. Massoulie, and C. Diot, "The diameter of opportunistic mobile networks," in *Proc. ACM CoNEXT*, 2007.
- [5] F. De Pellegrini, D. Miorandi, I. Carreras, and I. Chlamtac, "A graph-based model for disconnected ad hoc networks," in *Proc. IEEE INFOCOM*, 2007.
- [6] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *Proc. IEEE INFOCOM*, 2006.
- [7] V. Lenders, J. Wagner, S. Heimlicher, M. May, and B. Plattner, "An empirical study of the impact of mobility on link failures in an 802.11 ad hoc network," *Wireless Communications, IEEE*, vol. 15, no. 6, 2008.
- [8] H. Cai and D. Y. Eun, "Crossing over the bounded domain: From exponential to power-law inter-meeting time in MANET," in *Proc. ACM MobiCom*, 2007.
- [9] A. Vahdat and D. Becker, "Epidemic routing for partially-connected ad hoc networks," Tech. Rep., 2000.
- [10] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," *Computer Networks*, vol. 51, no. 10, pp. 2867 – 2891, 2007.
- [11] N. Eagle and A. Pentland, "Reality mining: Sensing complex social systems," *Personal and Ubiquitous Computing*, vol. 10, no. 4, 2006.
- [12] P.-U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. D. de Amorim, and J. Whitbeck, "The accordion phenomenon: Analysis, characterization, and impact on DTN routing," in *Proc. IEEE INFOCOM*, 2009.




Fig. 1. Influence of bundle size on delivery probability for different values of maximum delay (d). Each value of d corresponds to two lines: its upper and lower bounds.

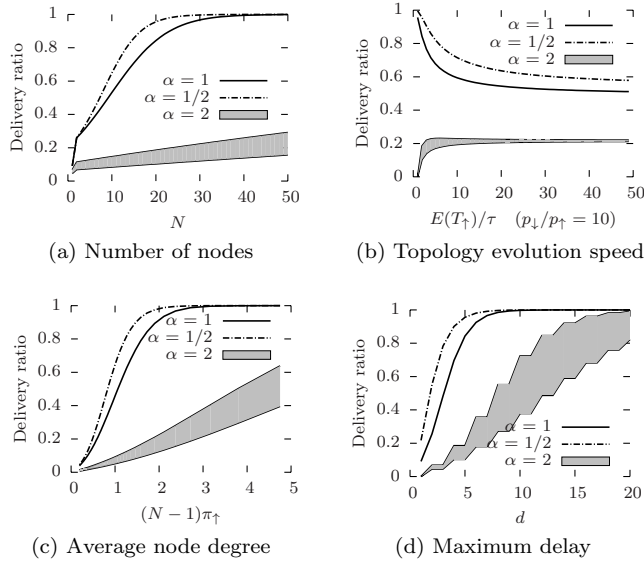


Fig. 2. Influence of model parameters on the delivery probability. When unspecified, $N = 20$, $p_{\downarrow} = 1/2$, $p_{\uparrow} = 1/20$, $d = 5$. Maximum delay and average link lifetime are expressed in number of time steps.




Fig. 3. Predicting delivery ratio for different bundle sizes in Rollernet with a 5-minute maximum delay.




Fig. 4. Rollernet: Delivery ratio vs. bundle size for various maximum delay values.